

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
23 August 2001 (23.08.2001)

PCT

(10) International Publication Number
WO 01/61478 A2

(51) International Patent Classification⁷: **G06F 9/38**

(74) Agent: **DEN BRABER, Gerard**; Internationaal Octrooibureau B.V., Prof Holstlaan 6, NL-5656 AA Eindhoven (NL).

(21) International Application Number: **PCT/EP01/00513**

(22) International Filing Date: 18 January 2001 (18.01.2001)

(81) Designated States (*national*): CN, JP, KR.

(25) Filing Language: English

(84) Designated States (*regional*): European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR).

(26) Publication Language: English

(30) Priority Data:
09/505,985 16 February 2000 (16.02.2000) US

Published:

— without international search report and to be republished upon receipt of that report

(71) Applicant: **KONINKLIJKE PHILIPS ELECTRONICS N.V.** [NL/NL]; Groenewoudseweg 1, NL-5621 BA Eindhoven (NL).

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(72) Inventor: **STRAVERS, Paul**; Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL).



WO 01/61478 A2

(54) Title: **SYSTEM AND METHOD FOR REDUCING WRITE TRAFFIC IN PROCESSORS**

(57) **Abstract:** An instruction execution device and method are disclosed for reducing register write traffic within a processor. The instruction execution device includes an instruction pipeline for producing a result for an instruction, a register file that includes at least one write port for storing the result, a bypass circuit for allowing access to the result, a means for indicating whether the result is used by only one other instruction, and a register file control for preventing the result from being stored in the write port when the result has been accessed via the bypass circuit and is used by only one other instruction.

System and method for reducing write traffic in processors

The present invention pertains generally to the field of digital computation circuits, and in particular, the invention relates to a system and method for an instruction execution device for use with a processor.

5

General motivational criteria exist for the design of microprocessors, for example, to reduce power consumption and size of such devices and as well reducing overall cost. In particular, one technological development in this regard has been the development of instruction execution architectures that implement a number of simultaneous parallel
10 instructions.

Systems and methods are known that provide instruction execution architectures of the type noted above, for example, microprocessor Instruction Set Architectures (ISAs). Typically, the implementation of such ISAs employs a so-called "pipeline" method to overlap different execution stages of subsequent instructions.

15

A conventional four-stage pipeline employs a (1) Fetch, (2) Decode, (3) Execute and (4) a Write-back. For data transfer type instructions such as a load instruction, one extra instruction pipeline stage is usually required.

20

In the first stage of the cycle, the processor fetches an instruction from memory. The address of the instruction to fetch is stored in the internal register, named the program counter, or PC. As the processor is waiting for the memory to respond with the instruction, it increments the PC. This means the fetch phase of the next cycle will fetch the instruction in the next sequential location in memory (unless the PC is modified by a later phase of the cycle).

25

In the decode phase, the processor stores the information returned by the memory in another internal register, known as the instruction register, or IR. The IR now holds a single machine instruction encoded as a binary number. The processor decodes the value in the IR in order to figure out which operations to perform in the next stage.

In the execution stage, the processor actually carries out the instruction. This step often requires further memory operations; for example, the instruction may direct the

processor to fetch two operands from memory (for example, storing them in operand registers), add them and store the result in a third location (the destination addresses of the operands and the result are also encoded as part of the instruction).

5 In the write-back stage of the pipeline, the result computed upstream in the pipeline is written (retired) to a destination register in a register file.

In another prior art pipeline method, circuitry is provided that allows operand or result values to bypass the register file. Using these bypass circuits, the operands or result values are already available to subsequent instructions before the operand-producing instructions are retired (e.g., written-back to register file).

10 There are, however, numerous shortcomings to these types of conventional pipelines. For example, conventional pipeline methods often require a large number of separate registers in a register file to adequately perform numerous simultaneous parallel instructions. The large register file typically contributes significantly to the overall power consumption. Moreover, each stage of the pipeline must be performed for each instruction
15 execution. These shortcomings, in turn, contribute to the power consumption and size of the processor. Accordingly, any decrease in the number of pipeline stages or circuit components needed for the pipeline required to perform instruction execution in a processor may (1) improve the over-all power consumption and (2) reduce the over-all size of the processor.

20

It is an object of the present invention to address some of the limitations of conventional instruction execution devices within processors, as discussed above.

It is a further object of the invention to provide a processor instruction set architecture that reduces the number of register file write operations in a pipeline method;
25 advantageously, in turn, reducing the overall power consumption of the processor. In addition, the size or silicon area of the register file is reduced due the decrease number of write ports being required.

The shortcomings associated with register write traffic in a processor are reduced or overcome by an arrangement in accordance with the principles of the present
30 invention in which an instruction execution device is provided. The instruction execution device includes an instruction pipeline for producing a result for an instruction, a register file that includes at least one write port for storing the result, a bypass circuit for allowing access to the result, a means for indicating whether the result is used by only one other instruction,

and a register file control for preventing the result from being stored in the write port when the result has been accessed via the bypass circuit and is used by only one other instruction.

In one embodiment of the present invention, the means for indicating whether the result is used by only one other instruction includes encoding each instruction. For
5 example, a so-called "dead value" field is designated in the "opcode" of each instruction to indicate whether the result will be used by only one other instruction.

In another embodiment of the invention, the means for indicating whether the first result is used by only the second instruction includes the instruction pipeline determining whether a result of an instruction in the instruction pipeline and another result of another
10 instruction in the instruction pipeline are designated for storage in the same write port in the register file. Since the write port in the register file is "reused" by a subsequent instruction already in the instruction pipeline, this is used to indicate that the first result will be used by only one other instruction.

These and other embodiments and aspects of the present invention are
15 exemplified in the following detailed disclosure.

The features and advantages of the present invention can be understood by reference to the detailed description of the preferred embodiments set forth below taken with
20 the drawings, in which:

Fig. 1 is a block diagram of one illustrative arrangement of an instruction execution device in accordance with the teachings of the present invention;

Fig. 2 is a block diagram of another illustrative arrangement of an instruction execution device in accordance with the teachings of the present invention; and

25 Fig. 3 is a flow chart depicting the process for reducing register write traffic in the arrangements of Figs. 1 and 2.

Some portions of the detailed descriptions that follow are presented in terms of
30 programs or algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired result. The steps are those requiring physical

manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical, magnetic or optical signals capable of being stored, transferred, combined, compared, and otherwise manipulated.

Referring now to Fig. 1, a block diagram is shown for illustrating one
5 embodiment of an instruction execution device in accordance with the teachings of the present invention. It will be recognized that Fig. 1 is simplified for explanation purposes and that the full processor environment suitable for use with the invention will comprise, for example, cache memory, RAM and ROM memory, compiler or assembler, I/O devices, etc., all of which need not be shown here. Generally, instruction execution device 10 uses an n-
10 stage pipeline instruction set register architecture (ISA) 12₁ through 12_n (hereinafter collectively known as "pipeline 12"), a conventional bypass circuit 14, a register file 16, and a register file control 18.

The pipeline 12 includes a number of pipeline stages (for example, Fetch, Decode, Execute and Write-back). It should be understood, however, that the invention is not
15 limited to a particular pipeline architecture. For example, the stages in a pipeline may include: instruction fetch, decode, operand fetch, ALU execution, memory access and write back of the operation results. Moreover, the chain of stages in the pipeline can be subdivided still more finely. The number of stages in the pipeline is an architectural feature, which can be changed according to the intended exploitation of instruction level parallelism.

20 Register file 16 includes at least one addressable destination write port 20 for storing data. As will be understood by persons skilled in the art, the register file can be any conventional database/indexing storage means that can store and allow access to records/data.

Register file control 18 contains the majority of logic, control, supervisory, translation functions required for controlling the operation of
25 writing-back write data to register file 16. Register file control 18 also includes programs for the operations functionally described in Fig. 3. As described in detail below, execution of these program implements the functionality necessary to reduce the number of registers file write operations in the pipeline.

Instructions can be classified as one of three major types: arithmetic/logic,
30 data transfer, and control. Arithmetic and logic instructions apply primitive functions of one or two arguments, for example addition, multiplication, or logical AND.

The timing of each stage depends on the internal construction of the processor and the complexity of the instructions it executes. The quantum time unit for measuring operations is known as a clock cycle. The logic that directs operations within a processor is

controlled by an external clock which, for example, may be a circuit that generates a square wave with a fixed period. The number of clock cycles required to carry out an operation determines the amount of time it will take.

I have observed that a large number of instructions produce a result value that is only consumed by a single instruction later in an algorithm, for example, an other instruction that enters the pipeline subsequently, (I will refer to the former instruction as the "producer" and to the latter as the "consumer"). Further, an analysis of the register data flow in the SPEC benchmarks revealed that 70% of all integer results and 80-99% of all floating point results are only consumed once, see an article written by M. Franklin et al., entitled "Register Traffic Analysis for Streamlining Inter-Operation Communication in Fine Grain Parallel Processors" published in The 25th Annual International Symposium on Microarchitecture (MICRO-25), pages 236-245, in December 1992. Furthermore, in a very large number of cases the consumer enters the pipeline architecture before the producer retires. This holds even stronger for specific types of processors, such as superscalar and VLIW processors. Accordingly, the consumer obtains the result value through a method other than the register file, for example, the bypass circuit. However, the result is nevertheless written-back to the register file.

More importantly, I have realized that the result value of the producer in such a situation does not have to be retired to the register file, since it will not be used by any other consumer. The result value of the producer effectively becomes a "dead value" as soon as it has been bypassed to the consumer. Thus, in accordance with one aspect of the invention, the register file control 16 determines whether a particular instruction should be written-back to register file 16 in stage-n, for example, the Write-Back stage, or not.

In the illustrative embodiment of Fig. 1, explicit encoding in each instruction is used to indicate whether the result value of a particular instruction will be used only one consumer (or by only other consumers in the pipeline). Preferably, a dedicated "dead value" bit in the instruction encoding (the so-called "opcode" of an instruction) is used, which is set or cleared by a compiler or assembler (not shown), depending on the degree of consumption. If the dead value bit is set, then the result value is not written-back to register file 16 in stage-n of pipeline 12. In particular, the dead value bit, as well as the instruction-id of the associated instruction, is provided to register file control 18, which in turn controls register file 16 via a write-enable signal to write-back the result value or not. Alternatively, if a dedicated bit is unavailable, then a few commonly used instructions can be selected (e.g.,

ADD and LOAD), which are assigned an alternative opcode to indicate the degree of consumption.

In another illustrative embodiment shown in Fig. 2, implicit encoding in each instruction is used to indicate whether the result value of a particular instruction will be used by only one consumer (or other consumers in the pipeline). Specifically, pipeline 12 determines whether another instruction in the pipeline, will use the same destination address (e.g., write port 20 in register file 16) as a subsequent instruction in the pipeline. This can be determined in a number of ways. For example, if the result value of the instruction in stage-3 12_3 is designated for the destination address of write port 20 in register file 16 and the result value of the instruction in stage-1 12_1 is also designated for the same destination address in register file 16, then the result value of the instruction in stage-3 is "alive" in stage-2, but is "dead" in stage-1 because its destination address will be reused by another instruction in the pipeline. In other words, only one consumer will use the result value. Alternatively, this method is also used when more than one consumer (in the pipeline) uses the result value. For example, if the result value of the instruction in stage-n 12_n is designated for the destination address of write port 20 in register file 16 and the result value of the (upstream) instruction in stage-1 12_1 is also designated for the same destination address in register file 16, then the result value of the instruction in stage-n is "alive" in stage-2 through stage n-1, but is "dead" in stage-1. In this manner, the result value can be used by numerous consumers while "alive". Accordingly, once a result value is "dead", it does not get written-back to register file 16.

In the illustrative embodiment of Fig. 2, conventional gate and logic circuits 22, for example OR gates, are used to provide a dead value bit, as shown in Fig. 2, to the register file control 18. The instruction-id of the associated instruction is also provided to register file control 18. However, as is well known in the art, in micro-architectures that support register renaming, gate and logic circuits are not required, since the processor already includes a mechanism to detect register reuse.

Fig. 3 is a flow chart showing the steps carried out within the arrangements Figs. 1 and 2 to implement a reduction of register write traffic. For simplicity, the embodiment of Fig. 1 is used to describe the operation of the principles of the present invention. However, it is to be understood that the steps described in Fig. 3 are equally applicable to the embodiment of Fig. 2.

With simultaneous reference to Figs. 1 and 3, the process contemplated by the invention is initiated in step 300 of Fig. 3, when an instruction enters pipeline 12. In step 302, the register file control 18 determines whether the dead value bit has been set, as described

above. If the dead value bit is not set, then the process continues in a conventional pipelining manner, represented by step 404. If the dead value bit is set, the register file control 18, in step 306, determines whether the result value of particular instruction has been already been used by an other instruction in pipeline 12, via the bypass circuit, for example. If the determination is negative, the process proceeds to step 304 and continues in a conventional pipelining manner. If the determination is affirmative, the process proceeds to step 308. In step 308, the register file control determines the result value corresponding to the received instruction id from pipeline 12 and discards it. For example, the register file control 18 does not allow the write-back stage to store the result value in a write port (or destination register) in register file 16.

The following merely illustrates the principles of the invention. It will thus be appreciated that those skilled in the art will be able to devise various arrangements which, although not explicitly described or shown herein, embody the principles of the invention and are included within its spirit and scope. Moreover, all statements herein reciting principles, aspects, and embodiments of the invention, as well as specific examples thereof, are intended to encompass both structural and functional equivalents thereof. Additionally, it is intended that such equivalents include both currently known equivalents as well as equivalents developed in the future, i.e., any elements developed that perform the same function, regardless of structure.

In the claims hereof any element expressed as a means for performing a specified function is intended to encompass any way of performing that function including, for example, a) a combination of circuit elements which performs that function or b) software in any form, including, therefore, firmware, microcode or the like, combined with appropriate circuitry for executing that software to perform the function. The invention as defined by such claims resides in the fact that the functionalities provided by the various recited means are combined and brought together in the manner which the claims call for. Applicant thus regards any means which can provide those functionalities as equivalent as those shown herein.

CLAIMS:

1. An instruction execution device for use in a processor, the instruction execution device comprising:
 - an instruction pipeline (12) for producing a first result for a first instruction;
 - a register file (16) connected to the instruction pipeline, the register file
 - 5 including at least a first write port (20) for storing the first result;
 - a bypass circuit (14) connected to the instruction pipeline (12), the bypass circuit (14) for allowing access to the first result;
 - means for indicating whether the first result is used by a second instruction in the instruction pipeline; and
 - 10 a register file control (18) connected to the instruction pipeline (12) and the register file (16), the register file control (18) for preventing the first result from being stored in the write port (20) when the first result has been accessed using the bypass circuit (18) and is used by the second instruction in the instruction pipeline (12).
- 15 2. The instruction execution device of claim 1, wherein the instruction pipeline (12) has a plurality of stages.
3. The instruction execution device of claim 2, wherein the instruction pipeline (12) processes a plurality of instructions in the plurality of stages.
- 20 4. The instruction execution device of claim 1, wherein the means for indicating whether the first result is used by only the second instruction includes encoding the first instruction to indicate whether the first result will be used by only the second instruction.
- 25 5. The instruction execution device of claim 4, wherein encoding the first instruction includes designating a field in the opcode of the first instruction indicate whether the first result will be used by only the second instruction.

6. The instruction execution device of claim 5, wherein each instruction is encoded to indicate whether its respective result is used by only one other instruction.

7. The instruction execution device of claim 3, wherein the means for indicating whether the first result is used by only the second instruction includes the instruction pipeline determining whether the first result of the first instruction in the instruction pipeline and a second result of the second instruction in the instruction pipeline are designated for storage in the first write port in the register file.

8. The instruction execution device of claim 7, wherein the means for indicating whether the first result is used by only the second instruction further includes gate and logic circuits (20).

9. A method of executing an instruction in a processor, the method comprising the step of:

generating a first result from a first instruction using an instruction pipeline;
allowing access to the first result using a bypass circuit (14);
determining whether the first result is used by a second instruction in the instruction pipeline;
preventing the first result from being stored in a register file (16) when the first result has been accessed and is used by only the second instruction.

10. The method of claim 9, wherein the determining step includes encoding the first instruction to indicate whether the first result will be used by only the second instruction.

11. The method of claim 9, wherein the determining step includes the instruction pipeline (12) determining whether the first result of the first instruction in the instruction pipeline (12) and a second result of the second instruction in the instruction pipeline (12) have the same storage address in the register file (16).

12. A computer-readable memory medium including code for processing an instruction, the code comprising:
instruction pipeline code for generating a first result from a first instruction;

access allowing code to the first result, while the first result in the instruction pipeline;

indicating code for identifying whether the first result is used by a second instruction in the instruction pipeline; and

5 preventing code to prevent storage of the first result when the first result has been accessed and is used by only the second instruction.

13. The memory medium according to claim 12, wherein the indicating code includes encoding the first instruction to indicate whether the first result will be used by only
10 the second instruction.

14. The memory medium according to claim 12, wherein indicating code includes the instruction pipeline determining whether the first result of the first instruction in the instruction pipeline and a second result of the second instruction in the instruction pipeline
15 have the same storage address in the register file.

1/3

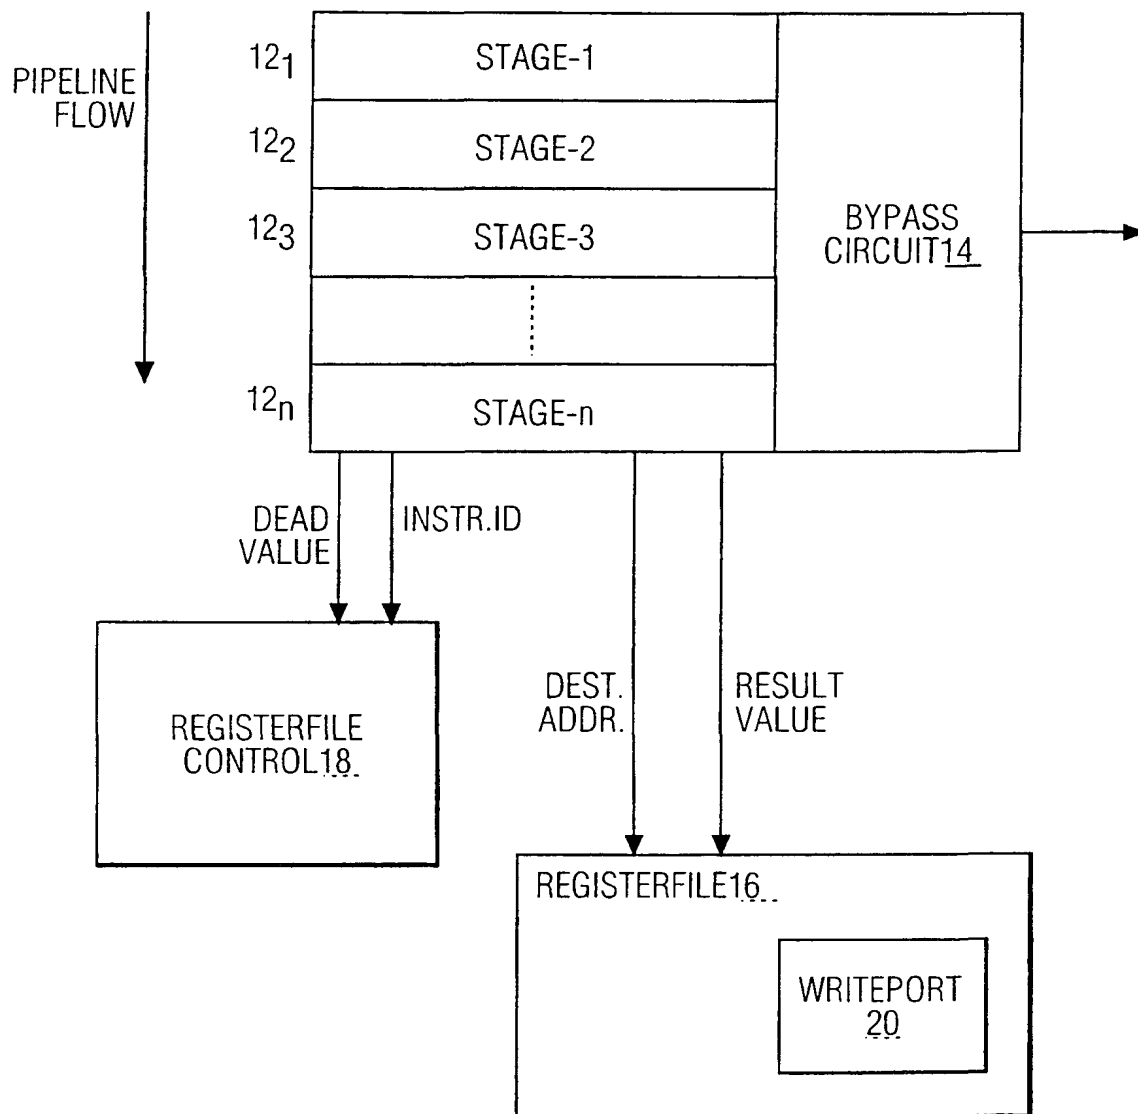


FIG. 1

2/3

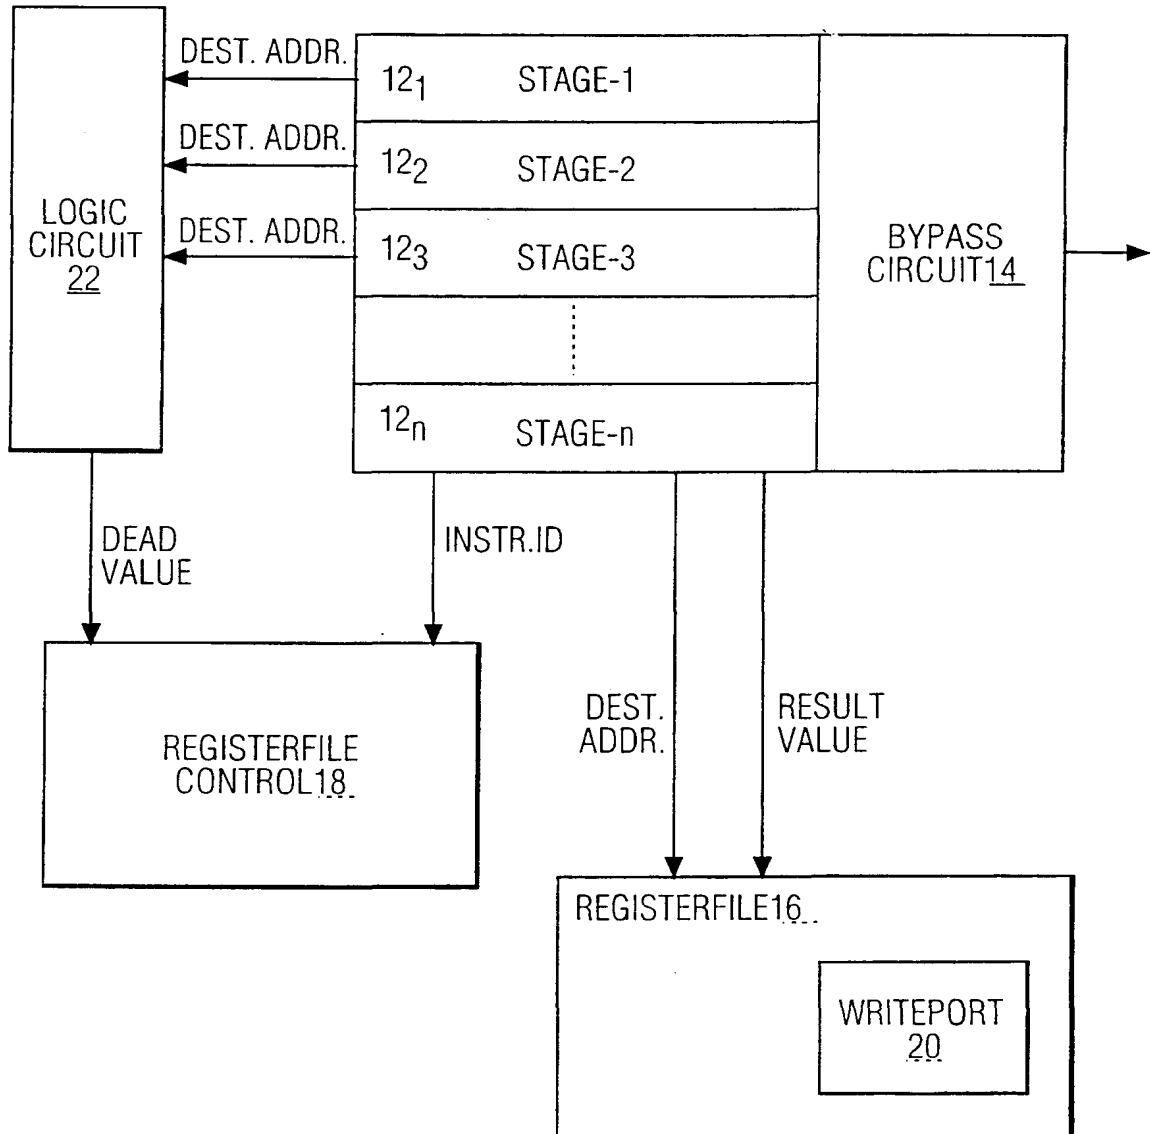


FIG. 2

3/3

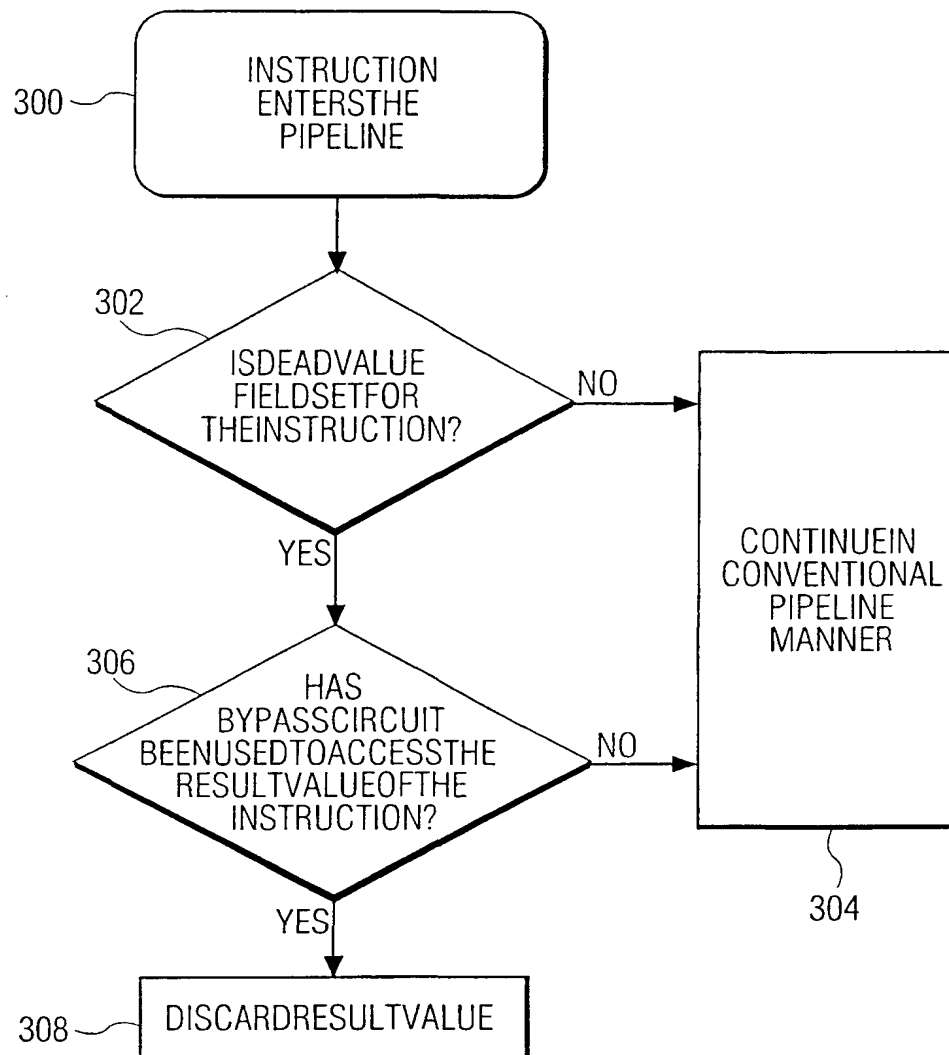


FIG. 3

(19) World Intellectual Property Organization
International Bureau



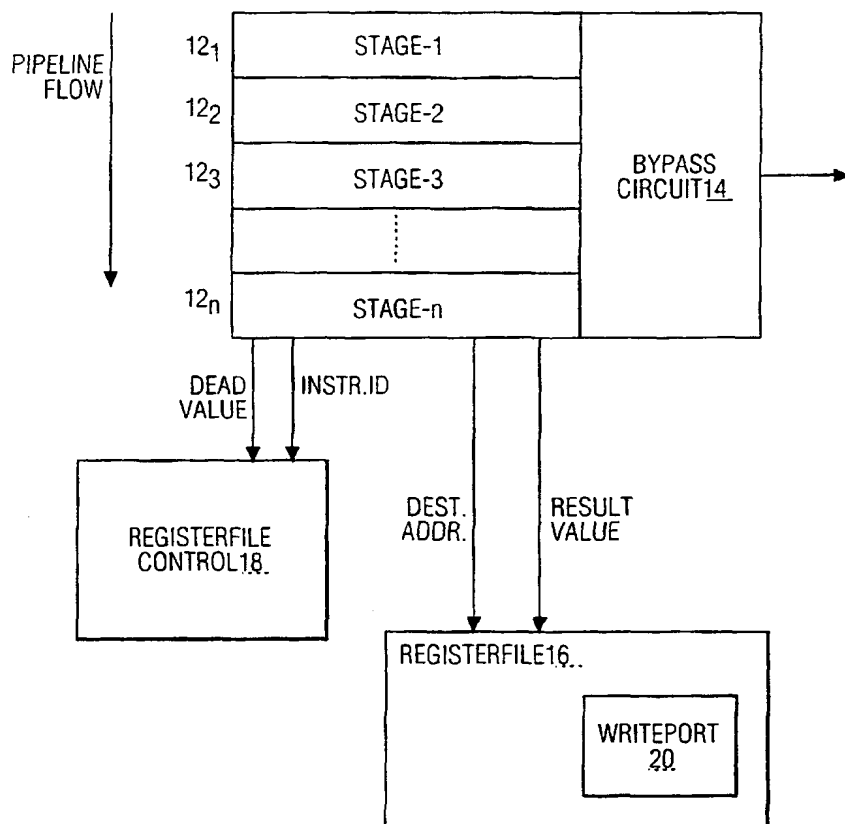
(43) International Publication Date
23 August 2001 (23.08.2001)

PCT

(10) International Publication Number
WO 01/61478 A3

- (51) International Patent Classification⁷: **G06F 9/38** (74) Agent: **DEN BRABER, Gerard**; Internationaal Octrooibureau B.V., Prof Holstlaan 6, NL-5656 AA Eindhoven (NL).
- (21) International Application Number: **PCT/EP01/00513**
- (22) International Filing Date: **18 January 2001 (18.01.2001)** (81) Designated States (*national*): **CN, JP, KR.**
- (25) Filing Language: **English** (84) Designated States (*regional*): **European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR).**
- (26) Publication Language: **English**
- (30) Priority Data: **09/505.985** **16 February 2000 (16.02.2000)** **US** Published: **— with international search report**
- (71) Applicant: **KONINKLIJKE PHILIPS ELECTRONICS N.V. [NL/NL]; Groenewoudseweg 1, NL-5621 BA Eindhoven (NL).** (88) Date of publication of the international search report: **27 December 2001**
- (72) Inventor: **STRAVERS, Paul**; Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL). *For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

(54) Title: **SYSTEM AND METHOD FOR REDUCING WRITE TRAFFIC IN PROCESSORS**



(57) Abstract: An instruction execution device and method are disclosed for reducing register write traffic within a processor. The instruction execution device includes an instruction pipeline for producing a result for an instruction, a register file that includes at least one write port for storing the result, a bypass circuit for allowing access to the result, a means for indicating whether the result is used by only one other instruction, and a register file control for preventing the result from being stored in the write port when the result has been accessed via the bypass circuit and is used by only one other instruction.

WO 01/61478 A3

INTERNATIONAL SEARCH REPORT

International Application No

PCT/EP 01/00513

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 G06F9/38

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, PAJ, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	LOZANO C L A ET AL: "EXPLOITING SHORT-LIVED VARIABLES IN SUPERSCALAR PROCESSORS" PROCEEDINGS OF THE ANNUAL INTERNATIONAL SYMPOSIUM ON MICROARCHITECTURE,US,LOS ALAMITOS, IEEE COMP. SOC. PRESS, vol: SYMP. 28, 29 November 1995 (1995-11-29), pages 292-302, XP000585373 ISBN: 0-8186-7349-4	1-3
A	the whole document ----- -/--	4-6,9, 10,12-14

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *Z* document member of the same patent family

Date of the actual completion of the international search

11 September 2001

Date of mailing of the international search report

18/09/2001

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040. Tx. 31 651 epo nl.
Fax: (+31-70) 340-3016

Authorized officer

Daskalakis, T

INTERNATIONAL SEARCH REPORT

International Application No

PCT/EP 01/00513

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>ARMIN LIEBCHEN ET AL: "DYNAMIC REORDERING OF HIGH LATENCY TRANSACTIONS USING A MODIFIED MICROPIPELINE"</p> <p>PROCEEDINGS OF THE INTERNATIONAL CONFERENCE ON COMPUTER DESIGN: VLSI IN COMPUTERS AND PROCESSORS, US, LOS ALAMITOS, IEEE COMP. SOC. PRESS, vol. -, 11 October 1992 (1992-10-11), pages 336-340, XP000344638</p> <p>ISBN: 0-8186-3110-4</p> <p>the whole document</p>	1
A	<p>US 4 228 497 A (GUPTA RAM K ET AL)</p> <p>14 October 1980 (1980-10-14)</p>	
A	<p>EP 0 497 485 A (IBM)</p> <p>5 August 1992 (1992-08-05)</p>	

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/EP 01/00513

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 4228497 A	14-10-1980	DE 2847575 A FR 2408174 A GB 2008819 A, B JP 1356447 C JP 54075935 A JP 61020013 B	23-05-1979 01-06-1979 06-06-1979 13-01-1987 18-06-1979 20-05-1986
EP 0497485 A	05-08-1992	US 5253349 A BR 9200241 A JP 5143323 A	12-10-1993 06-10-1992 11-06-1993

Form PCT/ISA/210 (patent family annex) (July 1992)